

# Supporting Named Entity Recognition and Document Classification in a Knowledge Management System for Applied Gaming

Philippe Tamla<sup>1</sup>, Florian Freund<sup>1</sup> and Matthias Hemmje<sup>2</sup>

<sup>1</sup>*Faculty of Multimedia and Computer Science, Hagen University, Germany*

<sup>2</sup>*Research Institute for Telekommunikation and Cooperation, Dortmund, Germany*

**Keywords:** Named Entity Recognition, Document Classification, Rule-based Expert System, Social Network, Knowledge Management System.

**Abstract:** In this research paper, we present a system for named entity recognition and automatic document classification in an innovative knowledge management system for Applied Gaming. The objective of this project is to facilitate the management of machine learning-based named entity recognition models, that can be used for both: extracting different types of named entities and classifying textual documents from heterogeneous knowledge sources on the Web. We present real-world use case scenarios and derive features for training and managing NER models with the Stanford NLP machine learning API. Then, the integration of our developed NER system with an expert rule-based system is presented, which allows an automatic classification of textual documents into different taxonomy categories available in the knowledge management system. Finally, we present the results of a qualitative evaluation that was conducted to optimize the system user interface and enable a suitable integration into the target system.

## 1 INTRODUCTION

The European research project Realizing and Applied Gaming Ecosystem (RAGE) is an innovative online portal and service-oriented platform for accessing and retrieving reusable software components and other related knowledge contents from the Web, such as research publications, source code repositories, issues, and online discussions. RAGE is used to support software reuse in the domain of applied gaming. Applied games (AG) or serious games (SG) aim at training, educating and motivating players, instead of pure entertainment (David R. and Sandra L., 2005). RAGE supports the integration with various social networks like Stack Exchange (“Hot questions”), or GitHub (“Build software better”). For instance, RAGE includes facilities to connect with the Stack Exchange REST API which enables an easy import of online discussions into its ecosystem. RAGE users can easily import multiple discussions from, for instance, the Stack Overflow social site, describe them with further meta information, classify them using an integrated taxonomy management system, and then finally retrieve useful information with faceted search that enables drilling down large set of documents.

Currently, the classification of text documents into

existing taxonomies in RAGE is done manually. The user has to, first, analyze the content of each document manually to understand the context in which this document is used. This is done by consulting the title and description of each imported document, as well as, analyzing all related meta-information (like keywords and tags), which are associated with this document. Once done, the user has to search for taxonomies that may be used to classify the imported document based on its content and metadata. This process can be very hard and requires the full attention of the user, because he or she needs to consult the document and taxonomy each time manually. With a large number of documents and multiple hierarchical taxonomies, it can be very time-consuming to classify documents in RAGE. To solve this problem, named entity recognition (NER) is generally applied because it can extract various knowledge contents (like named entities) from natural language texts (Nadeau and Sekine, 2007). The extracted knowledge content can then be used to automate the process of classifying text documents from various domains on the Web, using, for instance, an expert rule-based system.

NER has been widely used to recognize named entities in medical reports (K et al., 2017), news articles

(Newman et al., 2006), and software web documents (Ye et al., 2016; Tamla et al., 2019a). Techniques for NER vary from rule-based, over machine learning (ML), to hybrid methods. But, ML-based NER methods are more efficient on Web contents, because they include statistical models that can automatically recognize and classify named entities from very large and heterogeneous contents on the Web. The training of a machine learning-based NER model is however very challenging. It requires, besides very good programming knowledge, dealing with different technologies and pipelines for text analysis, natural language processing (NLP), machine learning and rule-based operations (Konkol, 2015). Errors in the initial stages of the pipeline can have snowballing effects on the pipeline's end performance. Therefore, facilitating the development, management, and execution of all necessary NER related tasks and pipelines will, not only reduce the effort to train new NER models but also contribute to optimizing the performance of the whole system. The goal of this research project is to develop and integrate a named entity recognition system into the RAGE ecosystem. The efficient integration of a NER system into the RAGE ecosystem will not only facilitate knowledge discovery (efficient extraction and analysis of named entities and their interrelationships), but also, enable an automatic classification of various text documents into existing taxonomies found in the ecosystem. After reviewing and comparing common systems and tools for named entity recognition and document classification, we present real-world use case scenarios and derive features for training and managing NER models with the Stanford NLP machine learning API. Then, the integration of our NER system together with the Drools expert rule-based system is presented, allowing an automatic classification of text documents into different taxonomy categories available in the knowledge management system. Finally, the results of a cognitive walkthrough, that served as a qualitative evaluation for optimizing the user interface and enabling a suitable system integration is shown.

## 2 STATE OF THE ART AND RELATED WORK

### 2.1 RAGE

As stated earlier, the RAGE social platform includes facilities for importing various text documents from the Web (like Stack Exchange discussions) into its ecosystem. These documents generally consist of a

title, a description and other metadata like tags, keywords, etc. An integrated *taxonomy management system* is also available for organizing and categorizing the textual materials into hierarchical taxonomies of the RAGE ecosystem. *Taxonomy* is the practice and science of classifying things and concepts including the principles underlining such classification (Sokal, 1963). It is used in RAGE to support faceted browsing, which is a technique allowing users to drill down their large number of search results, enabling faster information retrieval.

However, it is hard to classify documents with multiple taxonomies. The user can easily mix up one with another while analyzing and classifying a document into multiple hierarchical taxonomies. Each document (including its metadata like title, description, tags) have to be analyzed each time manually to be able to understand the context in which the document is used, before making a proper classification into the existing taxonomies. This process can be very challenging and time consuming, especially with multiple documents and various taxonomies having complex hierarchical structures. To fulfill the requirements of the project, a very desirable goal would be to develop and integrate a named entity recognition system into the RAGE system that can automatically recognize and classify various kinds of named entities from the multiple social networks connected with the ecosystem. Then, to apply an expert rule-based system that will enable an automatic document classification by reasoning about the extracted named entities, the hierarchical taxonomies and other textual features found in the existing RAGE documents.

### 2.2 Named Entity Recognition Techniques

NER techniques generally include handcrafted rules or statistical methods that rely on machine learning (ML) (Nadeau and Sekine, 2007), or even a combination of those. A NER technique is denoted as *rule-based* or *handcrafted* if all the parameters (including rules) that are used to identify and categorize named entities are defined manually by a human. *Machine learning* based techniques will use a computer to estimate those parameters automatically (Konkol, 2015). Existing ML techniques include *supervised learning* (parameter estimation is based on already annotated data), *semi-supervised learning* (parameter estimation uses only a small set of annotated data), and *unsupervised learning* (does not use annotated data for estimation). Most popular machine learning systems are relying on Conditional Random Fields (CRF), the state-of-the-art statistical modelling method for se-

quential text labelling (Sutton et al., 2007). CRF has been widely used with machine learning to support different NLP tasks, such as, *part-of-speech tagging* (Gimpel et al., 2010), *sentence splitting* (Tomanek et al., 2007) and *NER* (Ritter et al., 2011). Developing a machine learning-based NER system is however very challenges and requires a lot of data for model training. Often, *gazetteers* (dictionaries of specific named entities) are introduced as additional features to recognize unknown named entities - words that were not used in the training process. Likewise, *regular expressions* can be applied to optimize ML models, because they detect more complex named entities like compound words (Nagy et al., 2011).

Many factors can influence the performance of a NER system, such as a) The *language*. Some NER systems were developed for one specific language like English. b) The *named entity type*. For instance, the class of a datetime can be easily found if it only contains absolute dates (2003; 6.2.2005, April 5, 2011), but it can be difficult to detect relative dates (next Saturday, in December). c) The *domain* of the processed texts (corpora). If a classifier was trained using juristic texts, it will be difficult for this same classifier to deal with material originated from bioinformatics.

The standard measures for evaluating machine NER systems are *precision*, *recall* and *F1* for this task. Recall is the ratio of correct annotated NEs to the total number of correct NEs. Precision is the ratio of correct annotated NEs to the total number (correct and incorrect) of annotated NEs. F1 score is calculated from precision and recall and describes the balance between both measures. Most NER tools have functions to calculate precision, recall and F1 from a set of training and testing data.

### 2.2.1 Comparison of NER Tools

Many tools have been proposed in the literature for named entity recognition. We need to review and compare them to enable a suitable integration into our target system. Therefore, we introduce the following **selection criteria**: a) the chosen tool should not be limited to a specific type of text or knowledge domain b) should include a rich set of NLP features (including *NER*, *POS*, *Tokenization*, *Dependency Parsing*, *Sentiment Analysis*), c) must be stable, extendable, distributed as opensource, and should have an active community of developers.

Our solution is designed to classify a relatively small amount of data. The RAGE contents have a limited size and do not consist of many gigabytes of data. Therefore, we prefer to achieve good results with a high level of accuracy and do not need a very fast classification process which often results in lower

accuracy.

Our tool comparison is based on the work of Pinto (Pinto et al., 2016). According to our selection criteria, we exclude from our comparison non-opensource tools, tools without NER support, and those focusing only on specific data. To compare state-of-the-art tools, we added SpaCy, Spark NLP and Stanza to our list, because these tools arose in the last view years and may be relevant in our work. **GATE ANNIE**<sup>1</sup> is a more general solution for various NLP tasks. It was first developed to help software engineers and researchers working in NLP, but has been optimized to a more powerful system with an integrated user interface, which supports different data preprocessing tasks and pipeline executions. GATE is distributed with an integrated information extraction system called ANNIE that supports NER and many other NLP tasks. ANNIE relies on the JAPE<sup>2</sup> specification language, which provides finite state transduction over annotations based on regular expressions. Using the GATE interface, users can capture the provenance of machine- and human-generated annotated data to create new metrics for NLP tasks like named entity recognition. Additional metrics for more specific scenarios can be added, but this requires an existing implementation in the RAGE architecture, which introduces the overhead of familiarization with the entire GATE architecture. The **Natural Language Toolkit (NLTK)**<sup>3</sup> is a Python library that supports most of the common NLP tasks. It was launched in 2001 under the Apache license. Each NLP task is performed by an independent module and it is possible to train an own model for NER. The main disadvantage is that it lacks support for dependency parsing and an interface for the standard Universal Dependencies<sup>4</sup> dataset is missing. **Apache OpenNLP**<sup>5</sup> is written in Java and based on machine learning. Launched in 2004 and licensed under the Apache License, the software supports NER and many NLP tasks. But it lacks support for dependency parsing. The **Stanford CoreNLP**<sup>6</sup> is a Java-based tool suite from Stanford University that was launched in 2010. It supports all relevant NLP tasks, including NER and dependency parsing. CoreNLP can train new NER models independently from the data types, languages, or domain. Its API includes more than 24 different annotators for text annotation, regular expressions and language processing tasks. These annotators can be easily combined

<sup>1</sup><https://gate.ac.uk/ie/annie.html>

<sup>2</sup><https://gate.ac.uk/sale/tao/splitch8.html>

<sup>3</sup><https://www.nltk.org/>

<sup>4</sup><https://universaldependencies.org/>

<sup>5</sup><https://opennlp.apache.org/>

<sup>6</sup><https://stanfordnlp.github.io/CoreNLP/>

and executed sequentially in different pipelines. A REST service interface is also available, which can be used by other external systems for different NLP tasks execution. Thus, CoreNLP may be easily integrated with a rule-based expert system to support the automatic classification of documents in RAGE. Finally, the training of NER models is very flexible and customizable. CoreNLP includes nearly 100 parameters for CRF-based model training and performance fine-tuning, including other options for adding gazette lists that can recognize unknown named entities. CoreNLP is licensed under the GPLv3 and has a very big active community. Thus, state-of-the-art NLP methods and algorithms are permanently developed and integrated into the software. **Stanza**<sup>7</sup> is a Python Library, developed by Stanford University as a possible successor for CoreNLP. It was launched in 2019 under the Apache license. Even the system is rather new it supports many features needed in our work, only sentiment analysis is missing. The ML models trained by CoreNLP are not directly supported in Stanza and need to be trained again. Stanza brings a client to connect to the CoreNLP server, so it is possible to use CoreNLP features over this interface, which increases the complexity. **SpaCy**<sup>8</sup> is one of the newer systems for NLP that was launched in 2015. It is written in Python and was published under the MIT license. It is used to produce software for production usage, which should be easy to use and fast. SpaCy supports most of the common NLP features, including dependency parsing and features for training custom models for NER. But it lacks support for sentiment analysis. The main disadvantage for our purpose is, it focuses on fast classification, which leads to a lower accuracy compared to other systems. **Spark NLP**<sup>9</sup> is one of the most recent NLP tools that was released in 2017. It is a library build on top of Apache Spark and TensorFlow. It supports Python, Java and Scala and focuses the usage in production systems. It has more dependencies to get it up and running compared to other systems, due to the Apache Spark architecture. The supported NLP features include all relevant features, including dependency parsing and the training of a custom model for NER. Due to its young age, the community is not as big and active compared to others. On Stack Overflow are just a few questions with the tag “johnsnowlabs-spark-nlp” where “stanford-nlp” has over 3000. We decided to use the Stanford CoreNLP suite for our project. CoreNLP is the only NLP software which met all our requirements. The competitors may be

better or faster in one or another subtask, but overall CoreNLP seems to be the tool with the best mix of all required features. Especially the rich feature set in combination with an active and living community is a huge advantage of Stanford CoreNLP, compared to the other solutions.

## 2.3 Rule-based Expert Systems

Expert systems are rapidly growing technology of Artificial Intelligence (AI) that use human expert knowledge for complex problem-solving in fields like Health, science, engineering, business and weather forecasting (Koppich et al., 2009; Awan and Awais, 2011; Abu-Nasser, 2017). An expert system represents knowledge solicited by a human expert as data or production rules within a computer program (Abu-Nasser, 2017). These rules and data can be used to solve complex problems. For instance, a rule-based classification system can be applied to classify text documents into organized groups by applying a set of linguistic rules. The rules will instruct the system to use semantically relevant elements of the document and its contents to identify useful categories for automatic classification (Blosseville et al., 1992).

Over the last decades, many expert systems have been proposed but essentially all of them are expressed using IF THEN-like statements which contain two parts: the conditions and the actions. In the mathematical sense, a rule can be defined as  $X \implies Y$ , where  $X$  is the set of conditions (or antecedent) and  $Y$  is the set of actions (or the consequent). Rules are used to represent and manipulate knowledge in a declarative manner, while following the first-order logic in an unambiguous, human-readable form, and at the same time retaining machine interpretability. Rule-based systems generally include a “production memory” which contain a set of rules that are matched against facts stored in the “working memory” of an “inference engine” (Velickovski, 2016).

The **C Language Integrated Production System (CLIPS)** is a public domain software tool for building expert systems. It was developed by the NASA in 1985 (Velickovski, 2016). It has become one of the most used RBES in the market because of its efficiency and portability (Batista-Navarro et al., 2010). CLIPS was written C, and for C programming. But, it is now incorporating a complete object-oriented language for writing expert systems, called COOL. COOL combines the programming paradigms of procedural, object-oriented and logical languages. While CLIPS can separate the knowledge base (the expert rules) from its inference logic, it is not that user friendly in the formulation of rules like many other

<sup>7</sup><https://stanfordnlp.github.io/stanza/>

<sup>8</sup><https://spacy.io/>

<sup>9</sup><https://nlp.johnsnowlabs.com/>

systems (Velickovski, 2016).

Ten years after CLIPS, the **Java expert System Shell (JESS)** was launched by Ernest Friedman-Hill of Sandia National Lab (Velickovski, 2016) as a Java-based implementation of the CLIPS system. It supports the development of rule-based expert systems that can be tightly coupled to Java code and is often referred to as an expert system shell (Friedman-Hill, 1997). JESS is compatible with the CLIPS rule language, but a declarative language (called JessML) is also available for specifying rules in XML. JESS is free to use for educational and governmental purpose, but it is not an opensource software. There is no free source code under any available license<sup>10</sup>.

The **Drools** expert system is an opensource software that was first developed by Bob McWhiter (in 2001), and later on, absorbed by the JBoss organization (in 2005). Drools is based on Java and its rule definitions rely on IF...THEN statements which are easier to understand than the syntax provided by CLIPS and JESS. Drools rules can be also specified using a native XML format. The rule engine essentially is based on the Rete algorithm (Forgy, 1989), however, extended to support object-oriented programming in the rule formulation.

Drools is available under the Apache Software Foundation's opensource license. Because its easy and far more readable rule syntax, Drools has been widely used as an expert system in various domains (Cavalcanti et al., 2014). Therefore, we chose Drools to enable an automatic classification of text documents in the RAGE ecosystem.

### 3 SYSTEM DESIGN

Our system design relies on the user-centered design (UCD) approach by (Norman and Draper, 1986), which has proved to be very successful in the optimization of the product usefulness and usability (Vredenburg et al., 2002). Applying the UCD to design a system includes: a) understanding the context in which users may use the system, b) identifying and specifying the users' requirements, c) developing the design solutions, and finally, d) evaluating the design against users' context and requirements.

Our system allows any user (experts or novice developers) to customize and train a machine learning-based NER model in their domain of expertise. In the target system, the user starts with a named entity recognition definition, which is a set of parameters and configuration steps to train a named entity recog-

<sup>10</sup><https://jess.sandia.gov/jess/FAQ.shtml>

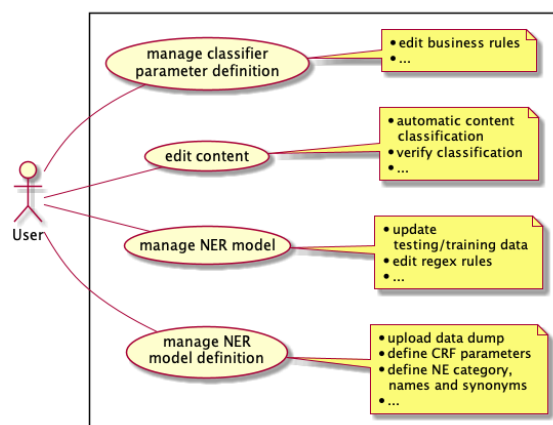


Figure 1: SNERC Use Case.

nition model using machine learning. With the support of the system, the user can upload a text corpus, define the named entity categories, and the named entity names (including their related synonyms) based on the requirements of the target domain. Then, he or she can customize all the conditional random fields and optimization parameters used to train a model using machine learning. The information about the NE categories, the NE names, and their related synonyms are used for the automatic annotation of the text corpus, using the BIO annotation mechanism which is integrated into our system. This is very useful because machine learning-based NER systems generally require a lot of annotated data for model training. However, while the system is able to suggest a first annotation of the text corpus, which can then be used for training and testing, it is necessary for the user to customize the testing data to avoid overfitting issues which may lead to very poor quality of the trained model (Konkol, 2015). Once a NER model is trained, the user can finally use it to construct flexible rules (by referring to the extracted named entities in the text) for automatic document classification in various domains. These rules are business rules and are constructed using a rule-based expert system. They will be used to represent and manipulate knowledge in a declarative manner using a set of WHEN... THEN statements in a human-readable form. The next sections will now provide an overview of relevant use cases and describe the overall architecture of the system.

#### 3.1 Use Case

Our use case diagram in figure 1 describes all tasks for a user to create a NER model definition, train a model, manage it, and finally use the trained model to support automated document classification in RAGE. We call

our system the Stanford Named Entity Recognition and Classification (SNER), as it relies on Stanford NLP for NER, and Drools for Document Classification. Our actor is a registered and logged-in user in KM-EP. There are four main actions:

“*manage NER model definition*” which includes:

- 1) uploading a data dump for use in the target domain,
- 2) defining the corresponding NE categories, names, and synonyms,
- 3) customizing CRF and performance parameters,
- 4) adding regular expressions to identify complex named entities (like Java 11.0),
- 5) preparing the NER model, which includes features for the automatic annotation of the text corpus and the splitting of the annotated text into testing and training data. Finally,
- 6) training the NER model using CronJobs and the Stanford NLP machine learning API.

“*manage NER model*” deals with the management of the created NER models. This includes reviewing the performance indicators like precision, recall and F1. The NER models can be edited or deleted. An upload of a pre-trained NER model is also possible.

“*manage classifier parameter definition*” deals with adding, editing, or deleting business rules that are used for classifying text documents into existing taxonomies. To create new rules, the user can select the taxonomies and NER models that are relevant for its specific domain.

The “*edit content*” action describes the steps, where a KM-EP content is edited and the automated classification suggestion is retrieved, supervised and saved.

### 3.2 System Requirements and Drools Extensions for Automatic Document Classification in RAGE

This section introduces our system requirements for document classification in RAGE. It also presents our feature extensions to create more complex rules based on Linguistic Analysis, Syntactic Pattern Matching, and Web Mining. Our features extensions will be implemented as a proof-of-concept for the classification of Stack Overflow discussions into RAGE taxonomies. Thus, we need to review the role of taxonomies in serious games to identify relevant taxonomies and validate our proof-of-concept. We can refer to our previous study about (Tamla et al., 2019b) to estimate which taxonomies may be relevant for the domain of serious games. In this research, we have applied the LDA statistical topic modelling to automatically discover 30 topics about serious games development, from which the following belong to the most popular ones: *Programming and Scripting Language, 3D-Modeling, Game Design, Rendering,*

*Game Engines, Game Physics, Networking, Platform, and Animation.*

We can now review the current state-of-the-art in taxonomies for serious games and select a list of taxonomies to be used in our proof-of-concept.

### 3.3 Serious Games-related Taxonomies

Taxonomies in serious games have many aspects and dimensions. Most relevant taxonomies for our work are related to 1) *Game genre*, 2) *programming languages*, 3) *video game tools*, 4) *machine learning algorithms*, and 5) *video game specification and implementation bugs*. Many researchers have proposed different hierarchical taxonomies in the domain of serious games. Their main objective was to elucidate the important characteristics of popular serious games and to provide a tool through which future research can examine their impact and ultimately contribute to their development (RATAN and Ritterfeld, 2009). Our first classification taxonomy reflects the *game genre* [GEN], as it is one the basic classification schemes proposed by researchers in the classification of serious games (RATAN and Ritterfeld, 2009; Buchanan et al., 2011; De Lope and Medina-Medina, 2017; Toftedahl and Henrik, 2019). A serious game can be classified based on the market [GEN/MAR](e.g. Education, HealthCare, Military), the game type [GEN/TYPE](board-game, card-game, simulation, role-playing game, toys, etc) or the platform [GEN/PLA]in which the game runs (Browser, Mobile, Console, PC) (RATAN and Ritterfeld, 2009). Many Stack Overflow discussions are already tagged with specific words like “education”, “board-game”, “simulation”, “console”. Therefore, we want to classify SG-related discussions in the game genre dimension. Second, our analysis of SG-related online discussions in Stack Overflow has revealed that developers of serious games are generally concerned with finding ways to implement new features using a specific programming language (or scripting) language [LANG]. So, a taxonomy in the programming language dimension is essential. To classify programming languages, we refer to Roy’s work (Van Roy et al., 2009) and use the *programming paradigm* as the main attribute in our work. We focus on serious game development, where existing game engines and tools for classic video game development are used, and we want to classify the Stack Overflow posts in this way. Third, (Toftedahl and Henrik, 2019) proposed a lightweight taxonomy to standardize the definition of common *tools, development environments*[TOOL/IDE], and game engines [TOOL/ENG] that are used for game development. We can use

this taxonomy as a classification scheme for the Stack Overflow posts. Fourth, another aspect is *machine learning [ML]*, the most trending aspect in serious games development. Machine learning is one of the main techniques used in reusable software components (Van der Vegt et al., 2016) and for creating intelligent learning systems. For instance, pedagogical systems use observational data to improve their adaptive ability, instead of relying on theoretical guidelines (Melo et al., 2018). This motivates us to integrate a machine learning-based classification scheme in our work. (Dasgupta and Nath, 2016) created such a scheme and gave a brief overview of state-of-the-art machine learning algorithms. We will use this in our work for classifying posts in the machine learning dimension. Our final dimension is regarding video game bugs [BUG]. As shown in our study, one of the main concerns of serious games developers (like most of the software developers) is to find solutions to fix their bugs, whether during the design or implementation of their games. (Lewis et al., 2010) developed in 2010 a taxonomy for video game bugs, which differentiate between specification bugs [BUG/SPEC] and implementation bugs [BUG/IMP]. A specification bug is generally referring to a wrong requirement in the game design document. This may refer to missing of critical information, conflicting requirements, or incorrectly stated requirements. A bug in an implementation is an error found in any asset (source code, art, level design, etc.) that is created to make the specification into a playable game (Varvaressos et al., 2017). A failure in an implementation is generally a deviation of the game's operation from the original game specification (Lewis et al., 2010).

### 3.4 Drools Extensions for Document Classification

This section presents our Drools extensions that is relevant to enable a flexible classification of text documents into the RAGE taxonomies. Our features extension rely on techniques for *Linguistic Analysis*, *Web Mining* and *Syntactic Pattern Matching*. Our classification system will be implemented as a standalone RESTful webservice so that it can be easily integrated within RAGE and any other external systems that may need to classify documents into predefined taxonomies.

#### 3.4.1 Linguistic Analysis

We use the Stanford NLP API to support linguistic analysis in our System. Stanford NLP supports many NLP tasks like part-of-speech tagging (POS),

tokenization, and NER. By analyzing specific part-of-speeches and recognizing various mentions of named entities discussion sentences, we can analyze the syntactic structure of each sentence. Then, we can refer to the **sentence components** (subject, predicate, object), the **sentence form** (whether it is *affirmative* (Liu et al., 2018) or *negative*), and the **sentence mood** (whether it is *interrogative* or *declarative*) to understand the structure of each sentence and derive its meaning. A similar approach was proposed by (Liu et al., 2018) for the classification of Stack Overflow discussions into software engineering-related facets, but this approach relied on hand-crafted rules for recognizing named entities in discussion posts. Instead of applying hand-crafted rules for NER, we will rely on our NER system to extract SG-related named entities (like game genres, programming languages, or game engines) from the existing text documents. To detect the *sentence form* and determine if a sentence is positive or negative, we will rely on the *StanfordNLP Sentiment Analysis API*<sup>11</sup>, as it includes a machine learning-based API for this purpose. We will rely on regular expressions to determine the *sentence mood*. We will consider a sentence to be *interrogative*, if it contains a question mark, or if it starts with an interrogative word (what, how, why, etc.) (e.g. what is the best way to record player's orientation?), otherwise the sentence is *declarative*. Using our linguistic analysis features, we can understand the meaning of each individual sentence, and use this information to derive the semantic of a document. Then, it becomes easier to group documents having similar semantic into a single taxonomy.

#### 3.4.2 Syntactic Pattern Matching

Research on web document mining has demonstrated that certain lexico-syntactic patterns matched in texts convey a specific relation (Liu and Chen-Chuan-Chang, 2004). Liu's study has revealed that many online questions belonging to similar topics have similar syntactic patterns. They found that many programming languages usually appear after a preposition, like **with** Java, **in** JavaScript. After carefully analyzing the title and description of some SG-related topics in Stack Overflow, we could easily observe similar behaviour for game genres, game engines and tools, such as **for** educational games, **in** Unity 3D, **with** GameMaker, etc. Thus, the categories of a question can be derived based on the syntactic patterns of its sentences.

Table 1 shows the list of our syntactic patterns that can be used to classify Stack Overflow discussions

<sup>11</sup><https://nlp.stanford.edu/sentiment/index.html>

Table 1: List of syntactic patterns.

Pattern	Description
PA	Entity or Term appears after a preposition
PB	Entity or Term appears before a preposition
SG	Entity or Term appears in the subject group
PG	Term appears in the predicate group
OG	Entity or Term appears in the object group
SA	The sentence is affirmative
SI	The sentence is interrogative
SP	The sentence is positive
SN	The sentence is negative
TT	Term combination $\langle term1 \rangle \langle term2 \rangle$ appears in a sentence
TTSG	Term combination $\langle term1 \rangle \langle term2 \rangle$ appears in the subject group
TTOB	Term combination $\langle term1 \rangle \langle term2 \rangle$ appears in the object group
TTPB	Term combination $\langle term1 \rangle \langle term2 \rangle$ appears before a preposition

into taxonomies of the RAGE system. Our syntactic pattern definition is based on a rich set of terms, term combinations, and standardized synonyms (Table 2), that we observed in various Stack Overflow discussions. Applying synonyms in our approach is very important to automatically detect name variations in text and enable a classification to perform better. For instance, we can use a pattern that includes the term “implement” and use the same pattern to identify texts that include the term “develop” or “build”. To achieve this goal, we will need to create a domain dictionary with a set of semantic classes, each of which includes a standardized term and its synonyms (Liu et al., 2018).

For each parameter in our defined template shown in Table 2, and for each taxonomy and category that the template applies to, we will use a list of popular terms found in Stack Overflow to instantiate our template and created a semantic class with each term. We will rely on the WordNet API<sup>12</sup> to create semantic classes of candidate synonyms using standardized terms. When a new term is added, all its synonyms should be identified using WordNet and then considered for inclusion. By combining different terms and synonyms, we can discover a wide range of expressions and term combinations and phrases used in the majority of SG-related discussions. For instance, the term combination  $\langle Best \rangle \langle Way \rangle$  can be used to

<sup>12</sup><https://wordnet.princeton.edu/>

identify posts containing the expressions: “best way“, “best strategy“, “proper design“, “optimal solution“, etc. This will allow us to have a more generic syntactic pattern definition that can easily scale in different domains compared to (Liu et al., 2018)’s system.

### 3.4.3 Document Structure Analysis

This feature is used to explore the structure of online text documents. We can refer to specific HTML elements to find out if a document contains a code snippets ( $\langle code \rangle \dots \langle /code \rangle$ ), bullet points ( $\langle ul \rangle \dots \langle /ul \rangle$ ), or even images ( $\langle img/ \rangle$ ). Exploring the structure of online discussion can help us to classify documents into specific taxonomies like *Programming Languages* or *Video Game Bugs*. A quality study of Stack Overflow online discussion (Nasehi et al., 2012) has revealed that explanations (generally represented using bullet points in the question bodies) accompanying code snippets are as important as the snippets themselves. Also, existing survey research on document structure analysis has demonstrated that analyzing the hierarchy of physical components of a web page can be very useful in indexing and retrieving the information contained in this document (Mao et al., 2003). For instance, if a Stack Overflow post, contains the word “bug” in its title, and one or more code snippets in its body, then it may be assigned to the *Implementation Category* of the *Video Game Bug Taxonomy*. Generally, such a discussion would include sentences like “How to **fix** my bug in...” or “How can I **solve** this issue... in my game” in its title or description body. Similarly, if a bug discussion includes terms like “requirement, design, or specification” in its title (e.g. I want to **fix** ... in my **specification**), with multiple bullet points in its description body, then it may indicate that the user is seeking help to solve an issue in a particular section of its design specification. In this case, the discussion post may be classified into the *Specification Bug* category of the *Video Game Bug Taxonomy*.

Our features extensions are very flexible and can be easily combined to construct even more complex rules in the Drools language. There is also no limitations for adding new extensions to classify text documents using our system.

## 3.5 System Architecture of SNERC

This section presents the system architecture of SNERC. Based on our use cases, we have defined 5 main components which will want to describe here. **NER Model Definition Manager** manages all the necessary definitions and parameters for model training using machine learning. It includes 3 main



Table 2: Template for synonym detection in Stack Overflow.

Taxonomy Category	Term	Term synonyms
Programming Language	< <i>implement</i> >	implement, develop, code, create, construct, build, set
Specification Bug	< <i>specify</i> >	design, require, define, determine, redefine
Implementation Bug	< <i>error</i> >	error, bug, defect, exception, warning, mistake
Specification Bug		
Game Engine	< <i>configure</i> >	configure, setup, adjust, adapt, optimize
-	< <i>howto</i> >	How to, How do (I,we), How can (I,we), How should (I,we)
... Bug	< <i>fix</i> >	fix, solve, remove, get rid of, eliminate

Table 3: Patterns for document structure analysis.

Pattern	Description
LS	Text contains multiple bullet points as HTML list
CS	Text contains one or multiple code snippets
IM	Text contains one or multiple images followed by a text description

classes. The first two, Named Entity Category and Named Entity, hold information about the domain-specific named entities names and categories. The third class, NERModelDefinition, is used to stored data like the model name, text corpus, gazette lists, and regex. We use the Stanford RegexNER API to construct and store complex rules, as they can easily be combined with already trained models.

**NER Model Trainer** is our second component that is used to prepare a NER model. This includes the automatic annotation of the domain text corpus (or data dump) based on the previously defined NE categories, NE names and synonyms. Our system is also able to split the annotated text corpus into testing and training data. The testing data, however, needs to be reviewed by a human expert and uploaded again to avoid overfitting, and thus a realistic calculation of precision, recall and F1 scores. When this is done, the NER Model Trainer component can execute the task for training a NER model using jobs and the Stanford CoreNLP. As the NER Model Trainer is written in Java and KM-EP is a PHP project, we designed it as a separate REST service component. This has further advantages. First, the service can be developed independently and does not affect KM-EP. Second, this service can be used separately from KM-EP as it is defined as a REST API. Other external systems will just need to define the input data in a JSON format and send them via an HTTP REST call to this service. The NER Model Trainer has a class called *NER Model Definition* which represents the corresponding GUI components in KM-EP. The Trainer class is used to control the training process.

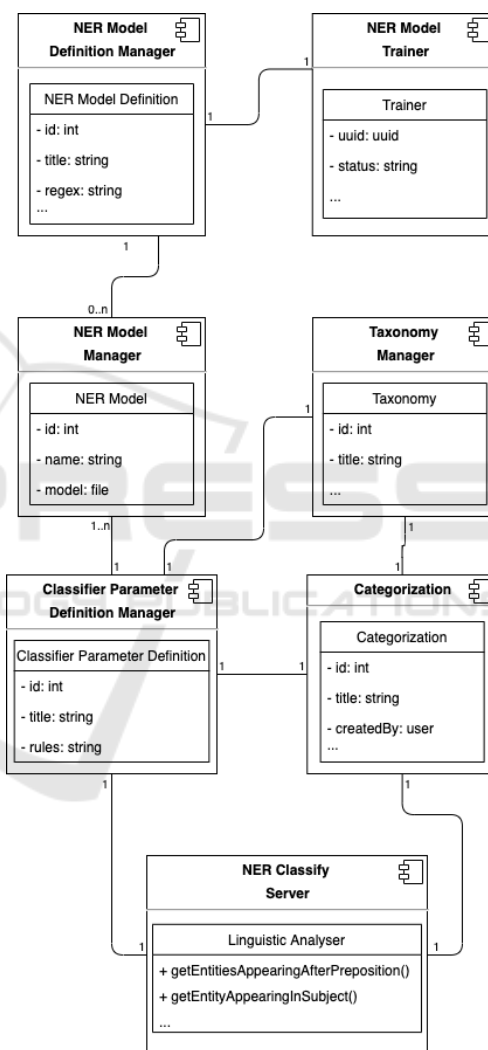


Figure 2: Model of the conceptual architecture.

**NER Model Manager.** This component is very straightforward since it only serves the storage of the trained NER models into the KM-EP filesystem so that they can be used by other systems like a linguistic analyzer or our document classification system. If a model is prepared with a NER Model Definition, users can update the created testing and training data

Table 4: Pattern matching rules for matching Stack Overflow discussion posts.

Pattern Matching	Taxonomy Categories	Examples
PA	LANG, GENRE, ...	< <b>How to</b> > to do animation with < <b>unity3d 5.2</b> >
(SG   OG) && SA	GENRE, ...	An < <b>Educational Game</b> > for learning prog. language.
(TT && SI)   PA	SPB	It might be an issue in the < <b>game</b> > < <b>design</b> > spec.
PB && CS	IMB	I am using a nstimer and it has a < <b>bug</b> > with my game loop < <b>code</b> >...< /code >

within the NER Model Manager to get better Precision, Recall and F1 scores. Also, the created Stanford Regex NER rules can be edited and updated. It is also possible to upload a StanfordNLP NER model that was trained with another system and use it in KM-EP. Figure 13 shows an example of a recognized named entity with the NER Model Manager.

**Classification Parameter Definition Manager.** This component is used to manage and store business rules in KM-EP. To construct business rules that mention named entities and can be used to classify documents into existing taxonomy categories, the design of the “Classification Parameter Definition Manager” component needs to include links to the “NER Model Manager”, “Content Manager” and “Taxonomy Manager” of KM-EP. We use the *Simple Knowledge Organization System (SKOS)* as the unique connection between our business rules and the taxonomy categories found in KM-EP. Even each taxonomy category in KM-EP has a SKOS persistent identifier representing the category.

**NER Classifier Server.** The NER Classify Server is our last component. It is developed as a standalone RestFul service to classify documents into taxonomies. To execute a document classification, the NER Classify Server needs information about the document (title, description, tags), the Drools rule, and references about the NER models, so that named entities can be used in the rule formulation. This information is sent to the server from KM-EP in a JSON format. With the provided document data and the references to the NER models, the server can now execute the NER, perform the synonym detection (with WordNet), and execute Linguistic Analysis, and Syntactic Pattern Matching on the Document structure and content. This analysis is done in the “classify()” method of a Java object, called Document. The analysis result is then stored into the properties of this object and can be used during the execution of Drools rules. The following code snippet shows the implementation of our Document.classify() method.

```
Server
  Document
    title
    description
    tags
```

```
...
classify()
  LinguisticAnalyzer.check(sentence)
    detectNamedEntities()
    detectSynonyms()
    appearsAfterPreposition()
    appearsBeforePreposition()
    isAffirmative()
    appearsInSubject()
    isSentencePositive()
  DocumentStructureAnalyzer(text)
    hasCodeSnippet()
    hasBulletPoint()
    hasImages()
```

### 3.5.1 System Service Implementation

To make the features of our implemented REST services available to the various KM-EP components, we created two new services in KM-EP. These services are used as an adaptor between KM-EP and its objects and our developed REST services. Each service bundles the features of the corresponding REST service and is connected with the KM-EP PHP API. The big advantage of relying on this service-based architecture is that, if we decide to change or update our REST APIs, we will only need to change the KM-EP services and leave their underline implementations untouched.

**NER Model Trainer Service.** The NER Model Trainer Service of KM-EP is used to connect with the NER Model Trainer REST service. As already discussed in the previous sections, this component includes the creation of a NER Model preview, the preparation of a NER Model and model training. Because the NER Models are created using the NER Model Trainer component, they need to be downloaded from there into KM-EP and deleted afterwards.

**Classifier Service.** The Classifier Service of KM-EP is used for the communication between KM-EP and the NER Classify Server REST service. To handle the automatic document classification, we first need to manage the NER Models using the NER Classify Server. Then, the Classifier Service of KM-EP can trigger the execution of the operation for adding or deleting NER Models by calling the NER Classify

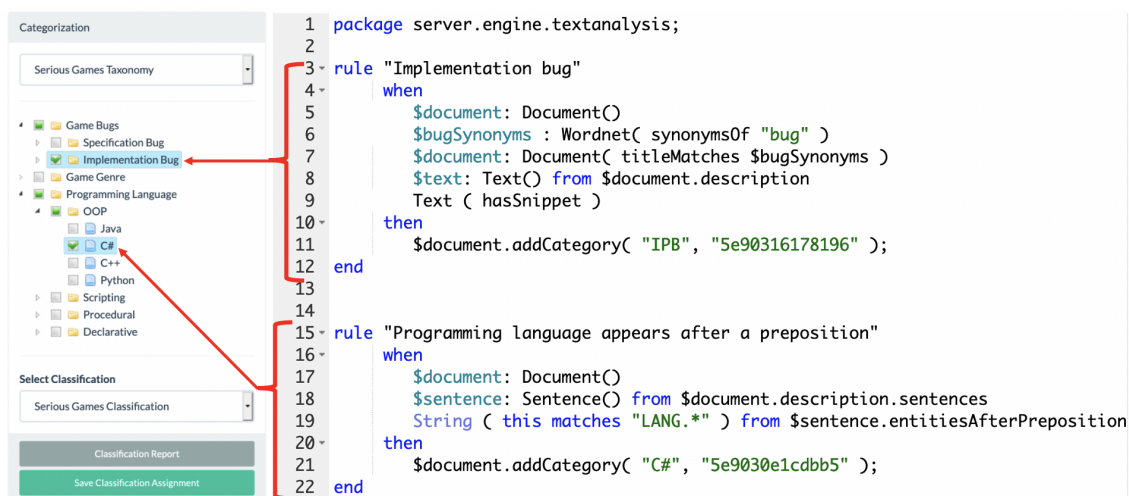


Figure 3: Selected categories and their rules.

Server. Furthermore, the Classifier Service will be able to trigger the automatic classification of documents to be suggested to the user.

### 3.6 Proof-of-Concept

After presenting our major use cases and showing details about our implemented components, we can now present a common use case scenario where Stack Overflow discussions about SG topics can be classified in RAGE.

With an existing NER model in the system, a classification parameter definition can be created with the Classification Parameter Definition Manager component to classify discussion texts into taxonomies of the system. For instance, there may be a Stack Overflow post like this in RAGE:

- Title: “bug in my game loop”
- Keywords: “cocoa-touch, nstimer”
- Description: “I am making a game on xcode 5. I am using a nstimer in C# and there may be a bug in my game loop. Can you help me please. All help is great. <code>...</code>”

According to our previous definition, we can create Drools rules to automatically classify this document into *Video Game Bug* and *Programming Language* taxonomies. First, we will start with the creation of a “Classification Parameter Definition”, where we select the desired taxonomy and NER models for named entity extraction. Then, we will construct our classification rules using the WHEN...THEN syntax provided by Drools. Based on the selected taxonomy, the NER models, and our rich set of features extensions, we can easily refer

to specific named entities (like C# (LANG), cocoa-touch (TOOL)) in our rule definitions and perform *Linguistic Analysis*, *Web Mining*, and *Syntactic Pattern Matching* based on the structure and content of our document. Figure 3 shows an example of such classification rules in the Drools language.

- Lines 6-7 (of rule 1) refer to our *WordNet* integration to detect if the term “bug” (or one of its synonyms) is included in the discussion title. Line 9 analyzes the document structure to identify if the post description includes a code snippet. Because both conditions are true, the document is automatically assigned to the *Implementation Bug* of the *Video Game Bug* taxonomy.
- Line 19 (of rule 2) checks the syntax of the post description to identify if a named entity of type LANG appears after a preposition. Since it is true, the post is assigned to the C# category of the *Programming Language* taxonomy.

To make it easier for the user to test the created rules, we implemented a form to test the developed rules. The user can input some text, execute the classification parameter definition and see a classification report with the results of the annotation and classification process. There is also a visualization of the NLP features detected by Stanford CoreNLP which is based on CoreNLP Brat<sup>13</sup>. The reports include the following information:

A list of persistent identifiers of the **detected categories**, an area for the detected **sentences** with the results of the Stanford CoreNLP features, representation of detected **Parts-of-Speech**, **detected NEs**,

<sup>13</sup><https://github.com/stanfordnlp/CoreNLP/tree/master/src/edu/stanford/nlp/pipeline/demo>

detected **basic dependencies** and the detected **sentiment**. For further analysis the original Stanford CoreNLP output is also available in JSON format in the GUI.

## 4 EVALUATION OF SNERC

After we implemented SNERC, it is needed to prove the usability of the system. There are several evaluation methods available to perform this task. Automated and formal methods are testing a system with a computer program, based on a formal specification, or with formal models. As it is difficult to create such a specification or model, we will not use one of these methods. Other methods like empirical methods involve a crowd of potential users of the system, which will perform common tasks in it. Such an evaluation is very resource-intensive and therefore not appropriate to our purpose. Informal methods are based on the knowledge and experience of the evaluating persons. It is known, that these methods create good results and detect many problems in a given system. On the other hand, they are not very difficult or expensive to implement, so they may be a good approach for our project. One of these informal inspection methods is the “Cognitive Walkthrough” (Polson et al., 1992), where a group of experts simulates a potential user of the system. The group navigates the system and tries to perform the typical steps to achieve the results a user tries to get. Potential problems and defects are documented and solved. Afterwards, the cognitive walkthrough may be repeated. We chose the cognitive walkthrough as an appropriate evaluation method for our system.

Our evaluation was performed in two steps. First, we performed a cognitive walkthrough in a collaborative meeting with three experienced experts: **Expert 1** is a very experienced professor and since many years Char of Area of Multimedia and Internet Application in the Department of Mathematics and Computer Science at FernUniversität in Hagen. **Expert 2** is a PhD, significantly responsible for the concept and design of KM-EP. **Expert 3** is a PhD student, researching in the area of serious games and named entity recognition.

First, the menu structure of SNERC was navigated exploratively, to simulate the navigation of a potential user in the system. Then each SNERC component was tested. Finally, the creation of an automated classification was evaluated. Within these steps, there were overall eight defects detected, which needed to be fixed. Then, a second evaluation was performed. We extended the expert group by two new evaluators: **Expert 4** is a PhD student, researching in the medical

area and emerging named entity recognition. **Expert 5** is a PhD student, researching in the area of advanced visual interfaces and artificial intelligence.

Within the second cognitive walkthrough all typical steps were performed, as a potential user would do it. There were no further defects detected. Expert 4 pointed to the problem of unrealistic performance indicators due to overfitting. This could be disproved with the possibility to supervise and edit the automatically generated testing data within the NER Model Manager. A further note was, SNERC may not be suitable to deal with huge data sets, because of its web-based GUI architecture. As KM-EP does not deal with such huge data sets this is not a real problem for our approach.

We saw the informal evaluation method lead to many results with a limited amount of time and resources. Nevertheless, an empirical evaluation with a bigger group of potential users should be done, to prove the usability and robustness of the system further.

## 5 CONCLUSIONS

In this research, we presented our integration of a named entity recognition and document classification tool into an innovative Knowledge Management System for Applied Gaming. Presentation of real-word use case scenarios for NER and automatic document classification has been highlighted and we saw, that it is possible to support users with the process of document classification through the use of named entity recognition in combination with a rule-based expert system. Our system has been successfully integrated and has been validated with a Cognitive Walkthrough. A future evaluation with a bigger group of potential users may help to gather further insights about the usage, usability and error handling of the system. Also, it should be analyzed if newer NLP tools evolved and may be suitable for our system. The use of standard tools for the management of the created ML models should also be taken into account.

## REFERENCES

- Abu-Nasser, B. (2017). Medical expert systems survey. *International Journal of Engineering and Information Systems (IJEAIS)*, 1(7):218–224.
- Awan, M. S. K. and Awais, M. M. (2011). Predicting weather events using fuzzy rule based system. *Applied Soft Computing*, 11(1):56–63.

- Batista-Navarro, R. T. B., Bandojo, D. A., Gatapia, M. A. J. K., Santos, R. N. C., Marcelo, A. B., Paganiban, L. C. R., and Naval, P. C. (2010). ESP: An expert system for poisoning diagnosis and management. *Informatics for Health and Social Care*, 35(2):53–63. Publisher: Taylor & Francis. eprint: <https://doi.org/10.3109/17538157.2010.490624>.
- Blosseville, M.-J., Hebrail, G., Monteil, M.-G., and Penot, N. (1992). Automatic document classification: natural language processing, statistical analysis, and expert system techniques used together. In *Proceedings of the 15th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 51–58.
- Buchanan, L., Wolanczyk, F., and Zinghini, F. (2011). Blending bloom’s taxonomy and serious game design. In *Proceedings of the International Conference on Security and Management (SAM)*, page 1. The Steering Committee of The World Congress in Computer Science, Computer . . . .
- Cavalcanti, Y. C., Machado, I. d. C., Neto, P. A. d. M. S., de Almeida, E. S., and Meira, S. R. d. L. (2014). Combining rule-based and information retrieval techniques to assign software change requests. In *Proceedings of the 29th ACM/IEEE international conference on Automated software engineering, ASE ’14*, pages 325–330. Association for Computing Machinery.
- Dasgupta, A. and Nath, A. (2016). Classification of machine learning algorithms. *International Journal of Innovative research in Advanced ngineering*, 3(3).
- David R. and Sandra L. (2005). *Serious games: Games that educate, train, and inform*. ACM.
- De Lope, R. P. and Medina-Medina, N. (2017). A comprehensive taxonomy for serious games. *Journal of Educational Computing Research*, 55(5):629–672.
- Forgy, C. L. (1989). Rete: A fast algorithm for the many pattern/many object pattern match problem. In *Readings in Artificial Intelligence and Databases*, pages 547–559. Elsevier.
- Friedman-Hill, E. J. (1997). Jess, the java expert system shell. Technical report, Sandia Labs., Livermore, CA (United States).
- Gimpel, K., Schneider, N., O’Connor, B., Das, D., Mills, D., and Eisenstein, J., editors (2010). *Part-of-speech tagging for twitter: Annotation, features, and experiments*.
- K, X., Zhou, Z., Hao, T., and Liu, W. (2017). A bidirectional lstm and conditional random fields approach to medical named entity recognition. *International Conference on Advanced Intelligent Systems and Informatics*.
- Konkol, I. M. (2015). *Named entity recognition*. Phd, PhD thesis, University of West Bohemia.
- Koppich, G., Yeng, M., and Ormond, L. (2009). Document management system rule-based automation. US Patent 7,532,340.
- Lewis, C., Whitehead, J., and Wardrip-Fruin, N. (2010). What went wrong: a taxonomy of video game bugs. In *Proceedings of the fifth international conference on the foundations of digital games*, pages 108–115.
- Liu, B. and Chen-Chuan-Chang, K. (2004). Special issue on web content mining. *Acm Sigkdd explorations newsletter*, 6(2):1–4.
- Liu, M., Peng, X., Jiang, Q., Marcus, A., Yang, J., and Zhao, W. (2018). Searching stackoverflow questions with multi-faceted categorization. In *Proceedings of the Tenth Asia-Pacific Symposium on Internetware*, pages 1–10.
- Mao, S., Rosenfeld, A., and Kanungo, T. (2003). Document structure analysis algorithms: a literature survey. In *Document Recognition and Retrieval X*, volume 5010, pages 197–207. International Society for Optics and Photonics.
- Melo, F., Mascarenhas, S., and Paiva, A. (2018). A tutorial on machine learning for interactive pedagogical systems. *International Journal of Serious Games*, 5(3):79–112.
- Nadeau, D. and Sekine, S. (2007). A survey of named entity recognition and classification. *Linguisticae Investigaciones*, 30:3–26.
- Nagy, I., Berend, G., and Vincze, V. (2011). Noun compound and named entity recognition and their usability in keyphrase extraction. In *Proceedings of the International Conference Recent Advances in Natural Language Processing 2011*, pages 162–169.
- Nasehi, S. M., Sillito, J., Maurer, F., and Burns, C., editors (2012). *What makes a good code example?: A study of programming Q&A in StackOverflow*. IEEE.
- Newman, D., Chemudugunta, C., Smyth, P., and Steyvers, M. (2006). *Analyzing entities and topics in news articles using statistical topic models*. International conference on intelligence and security informatics. Springer.
- Norman, D. A. and Draper, S. W. (1986). *User Centered System Design; New Perspectives on Human-Computer Interaction*. L. Erlbaum Associates Inc.
- Pinto, A., Gonçalo Oliveira, H., and Oliveira Alves, A. (2016). Comparing the Performance of Different NLP Toolkits in Formal and Social Media Text. In *5th Symposium on Languages, Applications and Technologies (SLATE’16)*, page 16 pages. Artwork Size: 16 pages Medium: application/pdf Publisher: Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik GmbH, Wadern/Saarbruecken, Germany.
- Polson, P. G., Lewis, C., Rieman, J., and Wharton, C. (1992). Cognitive walkthroughs: a method for theory-based evaluation of user interfaces. *International Journal of man-machine studies*, 36(5):741–773.
- RATAN, R. A. and Ritterfeld, U. (2009). Classifying serious games. In *Serious games*, pages 32–46. Routledge.
- Ritter, A., Clark, S., and Etzioni, O., editors (2011). *Named entity recognition in tweets: an experimental study*.
- Sokal, R. R. (1963). The principles and practice of numerical taxonomy. *Taxon*, pages 190–199.
- Sutton, C., McCallum, A., and Rohanimanesh, K. (2007). Dynamic conditional random fields: Factorized probabilistic models for labeling and segmenting sequence data. *Journal of Machine Learning Research*, 8:693–723.

- Tamla, P., Böhm, T., Hemmje, M., and Fuchs, M., editors (2019a). *Named Entity Recognition supporting Serious Games Development in Stack Overflow Social Content*. International Journal of Games Based Learning.
- Tamla, P., Böhm, T., Nawroth, C., Hemmje, M., and Fuchs, M., editors (2019b). *What do serious games developers search online? A study of GameDev Stack-exchange*, volume VOL. 2348 of *PROCEEDINGS OF THE 5TH COLLABORATIVE EUROPEAN RESEARCH CONFERENCE (CERC 2019)*. CEUR-WS.ORG.
- Toftedahl, M. and Henrik, E., editors (2019). *A Taxonomy of Game Engines and the Tools that Drive the Industry*, DIGRA INTERNATIONAL CONFERENCE 2019: GAME, PLAY AND THE EMERGING LUDO-MIX.
- Tomanek, K., Wermter, J., and Hahn, U. (2007). Sentence and token splitting based on conditional random fields. In *Proceedings of the 10th Conference of the Pacific Association for Computational Linguistics*, volume 49, page 57. Melbourne, Australia.
- Van der Vegt, W., Nyamsuren, E., and Westera, W. (2016). Reusable game software components and their integration into serious game engines. In *International Conference on Software Reuse*, pages 165–180. Springer.
- Van Roy, P. et al. (2009). Programming paradigms for dummies: What every programmer should know. *New computational paradigms for computer music*, 104:616–621.
- Varvaressos, S., Lavoie, K., Gaboury, S., and Hallé, S. (2017). Automated bug finding in video games: A case study for runtime monitoring. *Computers in Entertainment (CIE)*, 15(1):1–28.
- Velickovski, F. (2016). *Clinical decision support for screening, diagnosis and assessment of respiratory diseases: chronic obstructive pulmonary disease as a use case*. PhD thesis, University of Girona. Accepted: 2017-11-27T07:51:26Z Publisher: Universitat de Girona.
- Vredenburg, K., Mao, J.-Y., Smith, P. W., and Carey, T. (2002). A survey of user-centered design practice. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 471–478.
- Ye, D., Xing, Z., Foo, C. Y., Ang, Z. Q., Li, J., and Kapre, N., editors (2016). *Software-specific named entity recognition in software engineering social content*. IEEE.